A PCP Characterization of AM

Andrew Drucker
MIT

July 7, 2011

The class AM

- Arthur-Merlin (AM) protocols: a generalization of NP protocols **[Babai-Moran '88]**
- Explores the power of <u>randomness</u> in interaction with a prover.

The class AM

- $L \in \mathrm{AM}$ if there exists a polynomial-time algorithm $M(x, r, w)$, with
$$|r|, |w| \leq \mathrm{poly}(|x|) \ ,$$
such that:
    1. $x \in L \quad \Rightarrow \ \forall r \quad \exists w : \ M(x, r, w) = 1$;
    2. $x \notin L \quad \Rightarrow \ \Pr_r[\ \exists w : \ M(x, r, w) = 1] \ \leq \ 1/3$.
- $r =$ "random challenge"; $\quad w =$ "witness".

The class AM

- $(\Pi_Y, \Pi_N) \in \mathsf{AM}$ if there exists a polynomial-time algorithm $M(x, r, w)$, with
$$|r|, |w| \leq \mathrm{poly}(|x|) ,$$
such that:
  1. $x \in \Pi_Y \quad \Rightarrow \; \forall r \;\; \exists w : \;\; M(x, r, w) = 1$;
  2. $x \in \Pi_N \quad \Rightarrow \; Pr_r[\, \exists w : \;\; M(x, r, w) = 1] \; \leq \; 1/3$.

- $r =$ "random challenge"; $\quad w =$ "witness".

X

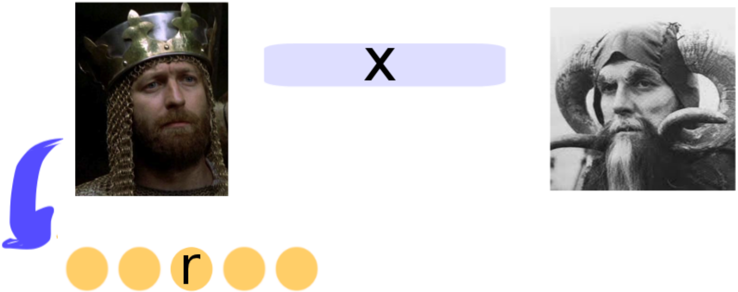
x
r

x

r

w

M(x, r, w) = 1 ---Accept!

x

r

w

AM vs. NP

- Clearly $\mathsf{AM} \supseteq \mathsf{NP}$.
- Is $\mathsf{AM} = \mathsf{NP}$? Is $\mathsf{AM} \subseteq \mathsf{NSUBEXP}$?

“Hardness vs. randomness”

- “Hardness vs. randomness” paradigm: sufficiently strong circuit lower bounds for exponential-time classes imply nontrivial derandomization of AM, even up to AM = NP.

  [Miltersen, Vinodchandran ’99; Shaltiel, Umans ’09]

- Gives a plausible reason to believe that AM = NP;

- But, not a currently viable approach to actually prove it!

“Hardness vs. randomness”

- Alternative approaches to AM vs NP?
- (**Caution:** any proof of AM = NP will imply some new circuit lower bounds;
  but <u>weaker</u> than those needed for hardness vs. randomness approach.)

Another approach

1. Identify "easiest" AM-hard problems;
2. Attack them with new algorithmic ideas.

- This work:

  gives candidate for (1), based on PCPs;

  shows obstacles to one algorithmic approach.

PCPs for AM

- We give a "PCP characterization of AM":
- For every $L \in \mathsf{AM}$, there's an AM protocol for L in which Arthur looks at only $O(1)$ bits of the witness string, <u>and</u> $O(1)$ bits of the random challenge!

X

X

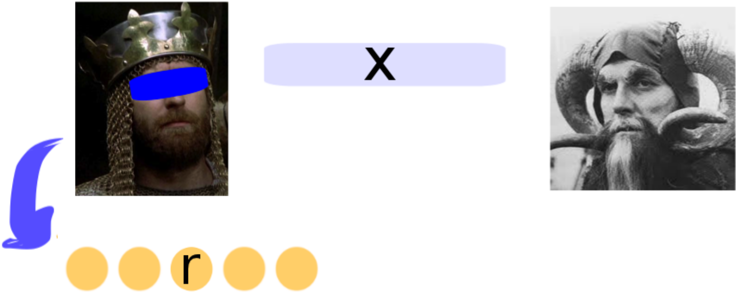
x
r

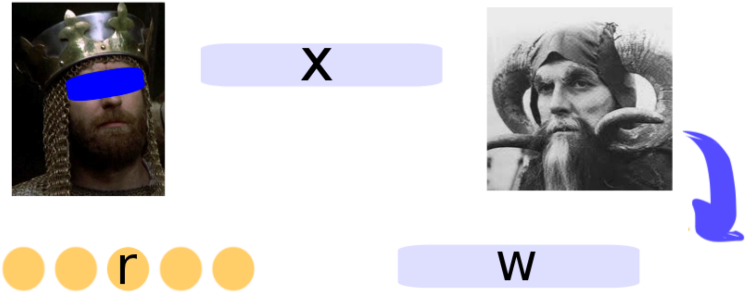
x
r
w

X

X


Accept!
x

Related work

- Idea of giving PCP-based complete problems for complexity classes other than NP is not new.
- Similar analogues of PCP Theorem given for:
  1. PH (the Polynomial Hierarchy)
     [Ko, Lin '94], [Haviv, Regev, Ta-Shma '07]
  2. PSPACE
     [Condon, Feigenbaum, Lund, Shor '95], [Drucker '11]
  3. IP = PSPACE
     [CFLS '97]

Switching views

- PCP Theorem can be described in terms of proof systems, or in terms of Constraint Satisfaction Problems (CSPs).
- Similarly with our result. We'll work with CSP viewpoint.

Stochastic CSPs

- k-CSPs: a family $\psi_1(z), \ldots, \psi_m(z)$ of constraints on variables z: each ψ_i is k-local.

  Let $\mathsf{Val}_\psi(z) =$ fraction of constraints ψ_i satisfied by z.

- Stochastic CSPs: $\psi(r, z)$

  $r =$ "random challenge" variables;

  $z =$ "witness/response" variables.

A complete problem for AM

- Say that $\psi(r,z)$ is risk-free if
$$\forall r \;\; \exists z : \; \mathrm{Val}_\psi(r,z) = 1 \; .$$
- Say that $\psi(r,z)$ is ε-risky if with probability $\geq 2/3$ over uniform r,
$$\forall z : \; \mathrm{Val}_\psi(r,z) < 1 - \varepsilon \; .$$

We show:

Theorem 1

There is an $\varepsilon > 0$ and a constant-size alphabet Σ such that, for stochastic 2-CSPs over Σ, it is AM-complete to distinguish between the cases

1. $\psi(r,z)$ *is risk-free;*
2. $\psi(r,z)$ *is ε-risky.*

Call this promise problem $\mathrm{Gap}-\mathrm{Stoch}-2\mathrm{CSP}_{\Sigma,\varepsilon}$.

Sketch of the proof

- Easy to see that the problem is in AM. Nontrivial direction: show it's AM-hard.
- Will show how to reduce any $L \in \mathsf{AM}$ to $\mathsf{Gap}-\mathsf{Stoch}-\mathsf{2CSP}_{\Sigma,\varepsilon}$.

  (Promise problems $\Pi \in \mathsf{AM}$ handled same way.)
- Given: an AM protocol $M(x, r, w)$ for a language $L \in \mathsf{AM}$.

Sketch of the proof

- Step 1: improve the soundness guarantee of M.

  Initial soundness $= 1/3$.

- Can drive down soundness to $(1/3)^{\mathrm{k}}$ by k-fold parallel repetition of M; but, blows up $|r|$ unacceptably.

- Instead, use randomness-efficient soundness amplification of **[Bellare, Goldreich, Goldwasser '93]**. Gives a new protocol M' for L, such that

$$x \notin L \quad \Longrightarrow \quad \Pr_r \left[\exists w : M'(x, r, w) = 1\right] \leq 2^{-\Omega(|r|)} .$$

Sketch of the proof

- Assume for simplicity that in M', we have $|r| \geq |w|$. (Can remove this assumption.)
- Let $C_x(r, w) := M'(x, r, w)$.

  C_x implementable by a $\mathrm{poly}(n)$-sized circuit.
- Then, rephrasing:

$$x \notin L \quad \Longrightarrow \quad \text{w.h.p. over } r \ ,$$

(r, w) is $\Omega(1)-$far in relative distance from $C_x^{-1}(1)$, for all w .

Sketch of the proof

- Step 2: Transform $C_x(r, w)$ into probabilistically checkable format.
  Key tool: *Prob. checkable proofs of proximity (PCPPs)*
  [Dinur, Reingold '04; Ben-Sasson et al. '04]

Theorem (Dinur '06)

There is a polytime transformation mapping a circuit $C(Y)$ to a 2-CSP $\psi(Y, z)$ over a constant-sized alphabet, such that for all y:

1. $C(y) = 1 \implies \exists z : \ \mathrm{Val}_\psi(y, z) = 1;$
2. *If y is δ-far from $C^{-1}(1)$, then $\forall z : \ \mathrm{Val}_\psi(y, z) < 1 - \Omega(\delta)$.*

Sketch of the proof

- Let $\psi_{C_x} = \psi_{C_x}(r, w, z)$ be the output of Dinur's reduction, applied to C_x.
  Let r be the random challenge vars; (w, z) witness-variables.

- Easy to check that $x \longmapsto \psi_{C_x}$ is the reduction we are looking for:

  1. $x \in L \implies \psi_{C_x}$ is risk-free;
  2. $x \notin L \implies \psi_{C_x}$ is $\Omega(1)$-risky.

- This proves $\mathrm{Gap} - \mathrm{Stoch} - \mathrm{CSP}_{\Sigma,\varepsilon}$ is AM-hard

  (for small $\varepsilon > 0$).

What next?

- Nontrivial derandomization for our AM-complete promise problem?

  ...haven't found one.

- How might we try?

What next?

- For a stochastic 2-CSP $\psi(r, z)$, what is the complexity of approximately optimizing over z, for randomly selected r?
- Perhaps easy, if we allow algorithm to depend nonuniformly on ψ...

A "randomized optimization" hypothesis

Hypothesis A

For any fixed $\delta, \varepsilon > 0$, and any stochastic 2-CSP $\psi(r, z)$ of size n, there is an "optimizer" circuit $\mathsf{OPT}_\psi(r)$, of size $\mathsf{poly}_{\delta,\varepsilon}(n)$ over r, such that with prob. $1 - \delta$,

$$\mathsf{Val}_\psi(r, \mathsf{OPT}_\psi(r)) \geq \max_z(\mathsf{Val}_\psi(r, z)) - \varepsilon .$$

A "randomized optimization" hypothesis

Claim

Hypothesis A implies AM = MA*.*

- Proof of Claim uses our AM-completeness result.
- If the optimizer circuits in Hyp. A can be NC^0 circuits, we'd get the stronger conclusion AM = NP.

Evidence against the hypothesis

- But—if NP is sufficiently hard, our plan fails:

Theorem 2
Suppose some $L \in \mathsf{NP}$ *is* $2/3$*-hard on average for circuits of size* $2^{\Omega(n)}$*.*

Then, Hyp A fails.

Proof sketch for Theorem 2

- Step 1: Our hardness assumption for $L \implies$

  $\exists$ a poly-time predicate $M(r, w)$ such that:

  1. $M(r, \cdot)$ is satisfiable w.h.p.; but,
  2. For any poly-sized circuit $C(r)$,

  $$\Pr_r[M(r, C(r)) = 1] \;=\; 2^{-\Omega(|r|)} \quad \text{(tiny)} \quad .$$

- Assume for simplicity: L balanced: $|L_n| = 2^{n-1}$.

Proof sketch for Theorem 2

- Natural idea for $M(r, w)$:

  r consists of many independent random strings of length n; $M(r, w)$ accepts iff w supplies proofs that at least a .49 fraction of them lie in L.

- $M(r, \cdot)$ is satisfiable w.h.p.—by Chernoff bounds!

- Any poly-sized circuit fails to satisfy $M(r, \cdot)$:

  Follows from hardness assumption on L and Direct Product theorems.

- Problem: uses too much randomness.

Proof sketch for Theorem 2

- Solution: use Impagliazzo-Wigderson PRG **[IW '97]**.
- To prove concentration property needed, apply recent Strong Chernoff Bound for Expander Walks **[Wigderson, Xiao '05], [WX '08], [Healy '08]**.

Proof sketch for Theorem 2

- Step 2: convert our predicate $M(r, w)$ into prob. checkable form.
- Idea: use PCPPs + error-correcting codes.
- This proves Theorem 2.

Summary

- Gave a new AM-complete problem, perhaps the "easiest" known.
- Advocated searching for an algorithmic attack on this problem to derandomize AM.
- Found obstacles to one natural approach.

Open Problems

- Complexity of $\mathrm{Gap}-\mathrm{Stoch}-2\mathrm{CSP}_{\Sigma,\varepsilon}$ when each "random" variable in $\psi(r,z)$ appears only $O(1)$ times in $\psi_1,\ldots,\psi_m$?
- Better hardness-vs-randomness results using $\mathrm{Gap}-\mathrm{Stoch}-2\mathrm{CSP}_{\Sigma,\varepsilon}$?
- New upper bounds on the power of AM protocols?